

# The Exchange Grid: a mechanism for data exchange between Earth System components on independent grids

V. Balaji<sup>a\*</sup>, Jeff Anderson<sup>b</sup>, Isaac Held<sup>c</sup>, Michael Winton<sup>c</sup>, Jeff Durachta<sup>c</sup>, Sergey Malyshev<sup>a</sup>, Ronald J. Stouffer<sup>c</sup>

<sup>a</sup>Princeton University  
PO Box 308, GFDL  
Princeton NJ 08542 USA

<sup>b</sup>NCAR Institute for Math Applied to Geophysics  
PO Box 3000  
Boulder CO 80307 USA

<sup>c</sup>NOAA Geophysical Fluid Dynamics Laboratory  
PO Box 308, GFDL  
Princeton NJ 08542 USA

We present a mechanism for exchange of quantities between components of a coupled Earth system model, where each component is independently discretized. The exchange grid is formed by overlaying two grids, such that each exchange grid cell has a unique parent cell on each of its antecedent grids. In Earth System models in particular, processes occurring near component surfaces require special surface boundary layer physical processes to be represented on the exchange grid. The exchange grid is thus more than just a stage in a sequence of regrid-ding between component grids.

We present the design and use of a 2-dimensional exchange grid on a horizontal planetary surface in the GFDL Flexible Modeling System (FMS), highlighting issues of parallelism and performance.

## 1. Introduction

In climate research, with the increased emphasis on detailed representation of individual physical processes governing the climate, the construction of a model has come to require large teams working in concert, with individual sub-groups each specializing in a different component of the climate system, such as the ocean circulation, the biosphere, land hydrology, radiative transfer and chemistry, and so on. The development of model code now requires teams to be able to contribute components to an overall coupled system, with no single kernel of researchers mastering the whole. This may be called the *distributed development process*, in contrast with the monolithic small-team process of earlier decades.

---

\*Corresponding author: balaji@princeton.edu

These developments entail a change in the programming framework used in the construction of complex Earth system models. The approach is to build code out of independent modular components, which can be assembled by either choosing a configuration of components suitable to the scientific task at hand, or else easily extended to such a configuration.

A specific issue that arises is how different components of the Earth system, say atmosphere and ocean, are discretized. Earlier generations of climate models used the *same* discretization, or simple integer refinement, for all components: thus, data exchange between components was a relatively simple point-to-point exchange. But any limitations on resolution of one component necessarily imposed itself on the other as well. Now it is increasingly common for each model component to make independent discretization choices appropriate to the particular physical component being modeled. In this case, how is, say a sea surface temperature from an ocean model, made available to an atmosphere model that will use it as a boundary condition? This is the *regridding problem*[6], subject to the following constraints when specialized to Earth system models:

- Quantities must be capable of being *globally conserved*: if there is a flux of a quantity across an interface, it must be passed conservatively from one component to the other. This consideration is less stringent when modeling weather or short-term (intraseasonal to interannual) climate variability, but very important in models of secular climate change, where integration times can be  $\mathcal{O}(10^6) - \mathcal{O}(10^8)$  timesteps.
- The numerics of the flux exchange must be stable, so that no limitation on the individual component timestep is imposed by the boundary flux computation itself.
- There must be no restrictions on the discretization of a component model. In particular, resolution or alignment of coordinate lines cannot be externally imposed. This also implies a requirement for *higher-order interpolation* schemes, as low-order schemes work poorly between grids with a highly skewed resolution ratio. Higher-order schemes may require that not only fluxes, but their higher-order spatial derivatives as well, be made available to regridding algorithms.

The independent discretization requirement extends to the time axis: component models may have independent timesteps. (We do have a current restriction that a coupling timestep be an integral multiple of any individual model timestep, and thus, timesteps of exchanging components may not be co-prime).

- The exchange must take place in a manner consistent with all physical processes occurring near the component surface. This requirement is highlighted because of the unique physical processes invoked near the planetary surface: in the atmospheric and oceanic boundary layers, as well as in sea ice and the land surface, both biosphere and hydrology.
- Finally, we require computational efficiency on parallel hardware: a solution that is not rate-limiting at the scalability limits of individual model components. Components may be scheduled serially or concurrently between coupling events.

The GFDL Flexible Modeling System (FMS)<sup>2</sup> [1] deploys software known as an *exchange grid* to solve this problem. We define an exchange grid in Section 2. We introduce the no-

---

<sup>2</sup><http://www.gfdl.noaa.gov/~fms>

tion of implicit coupling in Section 3, and demonstrate the coupling sequence used in FMS. In Section 4 we describe the parallelization of the exchange grid. Finally, in Section 5, we show basic performance characteristics of the FMS exchange grid at typical climate model resolutions. FMS and its exchange grid have been used in production for the models run at GFDL for the 2007 round of IPCC simulations (e.g [3]).

## 2. Definition of an exchange grid

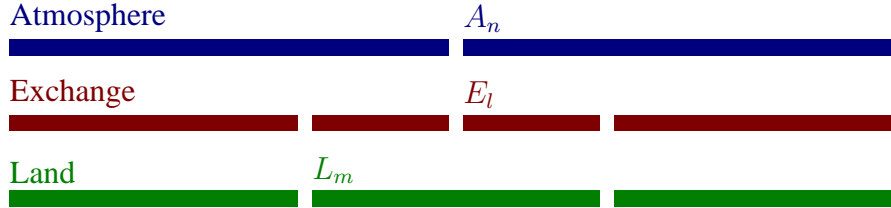


Figure 1. One-dimensional exchange grid.

A *grid* is defined as a set of *cells* created by *edges* joining pairs of *vertices* defined in a discretization. Given two grids, an *exchange grid* is the set of cells defined by the union of all the vertices of the two parent grids. This is illustrated in Figure 1 in 1D, with two parent grids (“atmosphere” and “land”). (Figure 2 shows an example of a 2D exchange grid, most often used in practice). As seen here, each exchange grid cell can be uniquely associated with exactly one cell on each parent grid, and *fractional areas* with respect to the parent grid cells. Quantities being transferred from one parent grid to the other are first interpolated onto the exchange grid using one set of fractional areas; and then averaged onto the receiving grid using the other set of fractional areas. If a particular moment of the exchanged quantity is required to be conserved, consistent moment-conserving interpolation and averaging functions of the fractional area may be employed. This may require not only the cell-average of the quantity (zeroth-order moment) but also higher-order moments to be transferred across the exchange grid.

Given  $N$  cells of one parent grid, and  $M$  cells of the other, the exchange grid is, in the limiting case in which every cell on one grid overlaps with every cell on the other, a matrix of size  $N \times M$ . In practice, however, very few cells overlap, and the exchange grid matrix is extremely sparse. In code, we typically treat the exchange grid cell array as a compact 1D array (thus shown in Figure 1 as  $E_l$  rather than  $E_{nm}$ ) with indices pointing back to the parent cells. Table 1 shows the characteristics of exchange grids at typical climate model resolutions. The first is the current GFDL model CM2 [3], and the second for a projected next-generation model still under development. As seen here, the exchange grids are extremely sparse.

The computation of the exchange grid itself could be time consuming, for parent grids on completely non-conformant curvilinear coordinates. In practice, this issue is often sidestepped by precomputing and storing the exchange grid. The issue must be revisited if either of the parent grids is adaptive.

The FMS implementation of exchange grids restricts itself to 2-dimensional grids on the

Atmosphere	Ocean	Xgrid	Density	Scalability
144×90	360×200	79644	$8.5 \times 10^{-5}$	0.29
288×180	1080×840	895390	$1.9 \times 10^{-5}$	0.56

Table 1

Exchange grid sizes for typical climate model grids. The first column shows the horizontal discretization of an atmospheric model at “typical” climate resolutions of  $2^\circ$  and  $1^\circ$  respectively. The “ocean” column shows the same for an ocean model, at  $1^\circ$  and  $\frac{1}{3}^\circ$ . The “Xgrid” column shows the number of points in the computed exchange grid, and the density relates that to the theoretical maximum number of exchange grid cells. The “scalability” column shows the load imbalance of the exchange grid relative to the overall model when it inherits its parallel decomposition from one of the parent grids.

planetary surface. However, there is nothing in the exchange grid concept that prevents its use in exchanges between grids varying in 3, or even 4 (including time) dimensions.

## 2.1. Masks

A complication arises when one of the surfaces is partitioned into *complementary components*: in Earth system models, a typical example is that of an ocean and land surface that together tile the area under the atmosphere. Conservative exchange between *three* components may then be required: crucial quantities like  $\text{CO}_2$  have reservoirs in all three media, with the *total* carbon inventory being conserved.

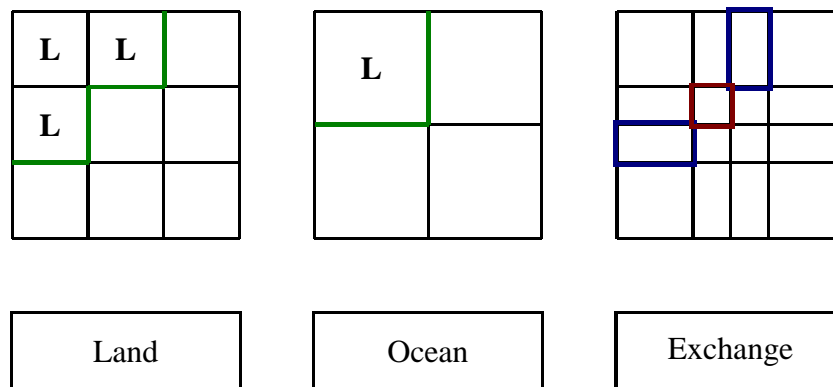


Figure 2. The mask problem. The land and atmosphere share the grid on the left, and their discretization of the land-sea mask is different from the ocean model, in the middle. The exchange grid, right, is where these may be reconciled: the red “orphan” cell is assigned (arbitrarily) to the land, and the land cell areas “clipped” to remove the doubly-owned blue cells.

Figure 2 shows such an instance, with an atmosphere-land grid and an ocean grid of different resolution. The green line in the first two frames shows the *land-sea mask* as discretized on

the two grids, with the cells marked **L** belonging to the land. Due to the differing resolution, certain exchange grid cells have ambiguous status: the two blue cells are claimed by both land and ocean, while the orphan red cell is claimed by neither.

This implies that the mask defining the boundary between complementary grids can only be accurately defined on the exchange grid: only there can it be guaranteed that the cell areas exactly tile the global domain. Cells of ambiguous status are resolved here, by adopting some ownership convention. For example, in the FMS exchange grid, we generally modify the land model as needed: the land grid cells are quite independent of each other and amenable to such transformations. We add cells to the land grid until there are no orphan “red” cells left on the exchange grid, then get rid of the “blue” cells by *clipping* the fractional areas on the land side.

## 2.2. Tile dynamics

A further complication arises when we consider tiles within parent grid cells. *Tiles* are a refinement within physical grid cells, where a quantity is partitioned among “bins” each owning a fraction of it. Tiles within a grid cell do not have independent *physical* locations, only their associated fraction. Examples include different vegetation types within a single land grid cell, which may have different temperature or moisture retention properties, or partitions of different ice thickness representing fractional ice coverage within a grid cell.

As the fractional quantity associated with each tile is different, we could associate an array dimensioned by tile with each exchange grid cell. The issue is that while the tile dimension may be large, the number of actual tiles with non-zero fraction is generally small. For instance, vegetation models often count 20 or more vegetation types (see e.g [2]); yet a single grid cell generally contains no more than two or three. It would therefore be inefficient to assign a tile dimension to the exchange grid: instead we model each tile as an independent cell and use the compaction as described above to eliminate tiles of zero fraction.

This implies that the compaction used to collapse the exchange grid sparse matrix is dynamic. We generally update the exchange grid every time a coupling event occurs. Compaction costs are negligible.

## 3. Implicit coupling

Fluxes at the surface often need to be treated using an implicit timestep. Vertical diffusion in an atmospheric model is generally treated implicitly, and stability is enhanced by computing the flux at the surface implicitly along with the diffusive fluxes in the interior. Simultaneously we must allow for the possibility that the surface layers in the land or sea ice have vanishingly small heat capacity. This feature is key in the design of the FMS coupler. Consider simple vertical diffusion of temperature in a coupled atmosphere-land system:

$$\frac{\partial T}{\partial t} = -K \frac{\partial^2 T}{\partial z^2} \quad (1)$$

$$\Rightarrow \frac{T_k^{n+1} - T_k^n}{\Delta t} = -K \frac{T_{k+1}^{n+1} + T_{k-1}^{n+1} - 2T_k^{n+1}}{\Delta z^2} \quad (2)$$

$$\Rightarrow \mathbf{A} \mathbf{T}^{n+1} = \mathbf{T}^n \quad (3)$$

This is a tridiagonal matrix inversion which can be solved relatively efficiently using an up-down sweep, as shown in Figure 3. The problem is that some of the layers are the atmosphere

and others are in the land. Moreover, if the components are on independent grids, the key flux computation at the surface, to which the whole calculation is exquisitely sensitive, is a physical process (e.g [7]) that must be modeled on the finest possible grid without averaging. Thus, the exchange grid, on which this computation is performed, emerges as an independent model component for modeling the surface boundary layer.

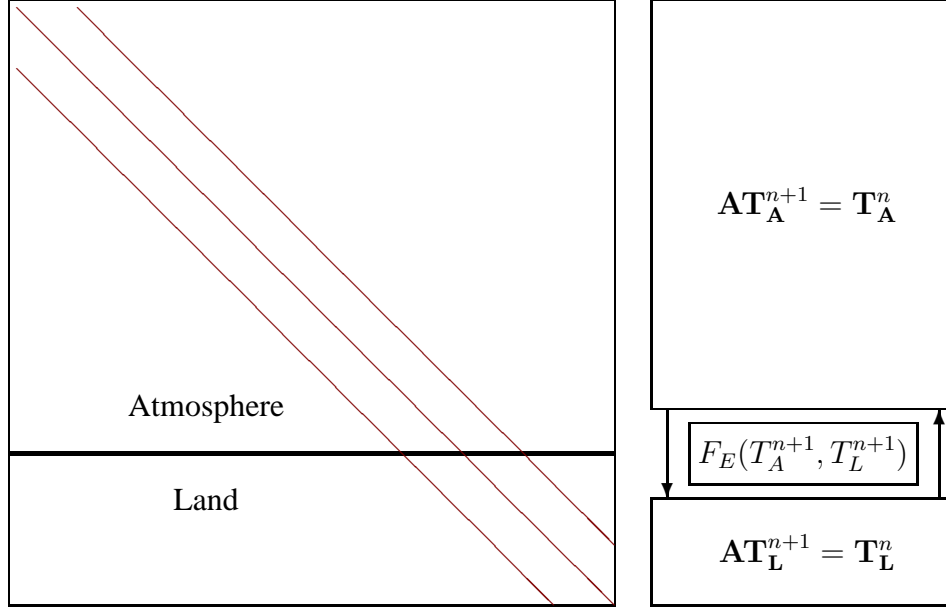


Figure 3. Tridiagonal inversion across multiple components and an exchange grid. The atmospheric and land temperatures  $T_A$  and  $T_L$  are part of an implicit diffusion equation, coupled by the implicit surface flux on the exchange grid,  $F_E(T_A^{n+1}, T_L^{n+1})$ .

The general procedure for solving vertical diffusion is thus split into separate up and down steps. Vertically diffused quantities are partially solved in the atmosphere (known in FMS as the “atmosphere\_down” step) and then handed off to the exchange grid, where fluxes are computed. The land or ocean surface models recover the values from the exchange grid and continue the diffusion calculation and return values to the exchange grid. The computation is then completed in the up-sweep of the atmosphere. Note that though we are computing vertical diffusion, some spurious horizontal mixing can occur as the result of regridding. More details are available in [4].

#### 4. Parallelization

Now we consider a further refinement, that of parallelization. In general, not only are the parent grids physically independent, they are also parallelized independently. Thus, for any exchange grid cell  $E_{nm}$ , the parent cells  $A_n$  and  $L_m$  (see Figure 1) may be on different processors. The question arises, to which processor do we assign  $E_{nm}$ ? The choices are,

1. to inherit the parallel decomposition from one of the parent grids (thereby eliminating communication for one of the data exchanges); or
2. to assign an independent decomposition to the exchange grid, which may provide better load balance.

In the FMS exchange grid design, we have chosen to inherit the decomposition from one side. Performance data (not shown here) indicate that the additional communication and synchronization costs entailed by choosing (2) are quite substantial, and of the same order as the computational cost as the flux computations on the exchange grid itself. Should the computational cost of the exchange grid grow appreciably, we may revisit this issue.

In choosing (1) we have also chosen to inherit the exchange grid parallelism from the side that uses dynamic tiling (generally, land and ocean surfaces; the atmosphere uses static tiling). The calls to the exchange grid are written to be entirely local, and all the communication is internal and on the static side. Note that when communication is on the side extracting data from the exchange grid onto the parent, we have a choice of performing on-processor sums *before* communication, which reduces data transfer volume. This results in loss of bit-reproducibility when changing processor counts, because of changes in summation order. We provide an option to send all data prior to summation as an option, when bitwise-exact results are needed (e.g for regression testing). This option adds about 10% to the cost at typical climate model resolutions.

## 5. Conclusions

We have described the design of a parallel exchange grid, used for conservative exchange of quantities between components of a coupled climate model, with sensitive flux computations in an intervening surface boundary layer. The physically most meaningful design is to perform these computations on the finest grid defined on the basis of the two parent grids. The exchange grid thus emerges as an independent model component, not just an intermediary in a regridding computation.

The exchange grid is designed not as a sparse matrix between all possible pairs of parent grid cells, but as a compact array pointing back to the parents. When parent grids have dynamic tiles, each tile appears as an independent cell on the exchange grid, rather than an extra tile dimension.

The grid may be precomputed and stored. If such grids are to be shared across a wide range of users, a standard grid specification must be accepted across the community. Efforts to develop a standard grid specification are underway. Masks defining the boundary between complementary grids are also best defined on an exchange grid.

Parallelization of the exchange grid is done on the basis of inheriting the decomposition from one of the parents, thus eliminating communication on one side of the exchange. The alternative, that of computing an independent decomposition for the exchange grid, was seen to be introducing much code complexity with the dubious promise of reward in the form of better load balance.

Current coupled climate computations in FMS [3] are carried out with a wide disparity in resolution ( $2^\circ \times 2.5^\circ$  in the atmosphere;  $1^\circ \times \frac{1}{3}^\circ$  in the ocean near the equator). The cost of the exchange grid component, including the flux computation, is about 12% at high scalability (180 processors). Much of the cost can be attributed to load imbalance: of a total of  $\sim 80000$

exchange grid cells,  $\sim 3000$  are on one processor in an extreme instance: thus limiting the scalability of the exchange grid relative to the overall model to about 0.29, as shown in Table 1.

The Flexible Modeling System website<sup>3</sup> contains the public-domain software for FMS, including its exchange grid. The site also contains links to documentation and papers on software features and design. The FMS exchange grid is a design prototype for the Earth System Modeling Framework [5], an emerging community standard for the construction of coupled Earth system models.

## Acknowledgments

V. Balaji is funded by the Cooperative Institute for Climate Science (CICS) under award number NA17RJ2612 from the National Oceanic and Atmospheric Administration, U.S. Department of Commerce. The statements, findings, conclusions, and recommendations are those of the author and do not necessarily reflect the views of the National Oceanic and Atmospheric Administration or the Department of Commerce.

## REFERENCES

1. V. Balaji. Parallel numerical kernels for climate models. In ECMWF, editor, *ECMWF Teracomputing Workshop*, pages 184–200. World Scientific Press, 2001.
2. Wolfgang Cramer, Alberte Bondeau, F. Ian Woodward, I. Colin Prentice, Richard A. Betts, Victor Brovkin, Peter M. Cox, Veronica Fisher, Jonathan A. Foley, Andrew D. Friend, Chris Kucharik, Mark R. Lomas, Navin Ramankutty, Stephen Sitch, Benjamin Smith, Andrew White, and Christine Young-Molling. Global response of terrestrial ecosystem structure and function to CO<sub>2</sub> and climate change: results from six dynamic global vegetation models. *Global Change Biology*, 7(4):357–373, April 2001.
3. T.L. Delworth, A.J. Broccoli, A. Rosati, R.J. Stouffer, V. Balaji, J.T. Beesley, W.F. Cooke, K.W. Dixon, J. Dunne, K.A. Dunne, J.W. Durachta, K.L. Findell, P. Ginoux, A. Gnanadesikan, C.T. Gordon, S.M. Griffies, R. Gudgel, M.J. Harrison, I.M. Held, R.S. Hemler, L.W. Horowitz, S.A. Klein, T.R. Knutson, P.J. Kushner, A.R. Langenhorst, H.-C. Lee, S.-J. Lin, J. Lu, S.L. Malyshev, P.C. D. Milly, V. Ramaswamy, J. Russell, M.D. Schwarzkopf, E. Shevliakova, J.J. Sirutis, M.J. Spelman, W.F. Stern, M. Winton, A.T. Wittenberg, B. Wyman, F. Zeng, and R. Zhang. GFDL's CM2 Global Coupled Climate Models. Part 1: Formulation and Simulation Characteristics. *Journal of Climate*, accepted for publication, 2005.
4. Isaac Held. Surface fluxes, implicit vertical diffusion, and the exchange grid. <http://www.gfdl.noaa.gov/fms/pdf/surf.pdf>, NOAA/GFDL, 1998. <http://www.gfdl.noaa.gov/fms/pdf/surf.pdf>.
5. Chris Hill, Cecelia DeLuca, V. Balaji, Max Suarez, Arlindo da Silva, and the ESMF Joint Specification Team. The Architecture of the Earth System Modeling Framework. *Computing in Science and Engineering*, 6(1):1–6, January/February 2004.
6. P. Knupp and S. Steinberg. *Fundamentals of Grid Generation*. CRC Press, 1993.
7. A.S. Monin and A.M. Obukhov. Basic laws of turbulent mixing in the ground layer of the atmosphere. *Tr. Geofiz. Inst. Akad. Nauk. SSSR*, 151:163–187, 1954.

---

<sup>3</sup><http://www.gfdl.noaa.gov/~fms>